

Specimen — illustrative example

Solstice Grid Services is a fictional organisation. This document comes from the scenario described on riskwoven.com/examples, and is not an assessment of any real entity. It is advisory output to support your own decisions: not legal advice, not a certification, and not an audit.

Written and reviewed by RiskWoven.com. Written by RiskWoven.com's AI drafting agent and checked by RiskWoven.com's independent review stage before release. The review is automated. It is not a human, legal or audit review.

AI Governance and Accountability Policy

Solstice Grid Services

Document type	policy	Version	1.0 Draft
Owner	Head of Engineering	Customer decision role	Chief Technology Officer
Effective date	Upon customer release	Review cycle	Annual
Framework basis	RiskWoven AI practical best practices		
Jurisdiction	Ontario, Canada		
Prepared for	Solstice Grid Services	Industry	Energy and utilities
Organization size	101-500	Locations	Not specified

Advisory draft for accountable-owner and professional review. Validate legal, regulatory, contractual, technical and operational requirements before customer release. This is not legal advice, a certification, or a compliance determination.

1. Purpose

This policy establishes the accountability framework for artificial intelligence systems used by Solstice Grid Services. It mandates that every AI system is inventoried, risk-assessed,

assigned an accountable owner, subject to proportionate human oversight, monitored throughout its operational life, and retired securely when no longer required.

The policy exists because Solstice Grid Services operates autonomous agents that act on grid telemetry and draft customer correspondence. These systems affect grid reliability, customer service and the handling of personal information for approximately 85,000 metered accounts. Governance failures in this context carry consequences for safety, regulatory standing and public trust that justify dedicated oversight.

2. Scope

This policy applies to:

- Every AI system developed, procured, hosted or operated by Solstice Grid Services, whether the system runs on Solstice Grid Services infrastructure or is provided by a third party.
- Every employee, contractor and third party who proposes, develops, deploys, operates, monitors or retires an AI system on behalf of Solstice Grid Services.
- AI capabilities embedded in larger platforms, including the agent orchestration platform, the model provider, the telemetry pipeline vendor and the mapping data provider.

The policy covers the full lifecycle: proposal, risk assessment, review, deployment, operation, material change, and retirement.

3. Definitions

AI system: Any software that uses machine learning, statistical inference, rule-based reasoning or generative modelling to produce outputs—predictions, recommendations, decisions, content or actions—that would otherwise require human judgement. This includes autonomous agents that read inputs and take actions, models that draft correspondence, and any component that learns or adapts from data.

System owner: The individual accountable for an AI system throughout its lifecycle. The system owner ensures the system is inventoried, risk-assessed, operated within its defined boundaries, monitored, and retired securely. Accountability cannot be delegated; tasks may be.

Risk tier: A classification that reflects the potential impact of an AI system if it fails, behaves unexpectedly or is misused. Solstice Grid Services uses three tiers—High, Medium and Low—determined by the method in Section 7.

Impact assessment: A documented evaluation of an AI system's potential effects on grid safety and reliability, on customers and employees, on personal information, and on regulatory obligations. The assessment identifies harms, their likelihood, the controls that reduce them, and the residual risk that remains.

Material change: Any modification to an AI system, its data inputs, its action boundaries, its hosting environment or its supplier arrangements that could alter its risk profile. Examples include: a new model version from a supplier, an expanded action boundary for an agent, a change in the data the system receives, a new subprocessor, or a change in the jurisdiction where data is processed.

4. Policy Statements

4.1 Solstice Grid Services maintains a complete, current inventory of every AI system within scope. No AI system may operate without an inventory entry.

4.2 Every AI system has a named system owner who is accountable for its compliance with this policy.

4.3 Every AI system undergoes an impact assessment before deployment and before any material change. The depth of assessment is proportionate to the system's risk tier.

4.4 No AI system may be deployed or materially changed without review by the appropriate authority for its risk tier.

4.5 Human oversight is assigned to every AI system. The form and intensity of oversight is proportionate to the system's potential impact.

4.6 Autonomous agents that act without human confirmation of each action operate only within a defined, tested action boundary. Any action outside that boundary halts the agent until a human clears it.

4.7 AI systems are monitored for performance degradation, drift, incidents, complaints and supplier changes throughout their operational life.

4.8 AI systems are retired securely when no longer required, with evidence preserved as mandated by the Records Retention and Disposal Policy.

4.9 Exceptions to this policy are time-limited, recorded, and accepted only by the authority specified for the relevant risk tier.

5. Roles and Accountabilities

Chief Technology Officer:

- Holds final review authority for High-risk AI systems and for any agent granted action capability on critical suppliers.
- Accepts residual High risk, for a fixed period, with a recorded review date.
- Resolves escalations that cannot be resolved at lower levels.
- Reviews this policy annually or following a significant incident.

Head of Engineering:

- Owns this policy and is accountable for its maintenance and communication.
- Holds review authority for Medium-risk AI systems.
- Maintains the AI system inventory and ensures entries are current.
- Escalates unresolved risks or disputes to the Chief Technology Officer.

System Owner (assigned per system):

- Accountable for the AI system throughout its lifecycle.
- Ensures the system is inventoried, assessed, operated within its boundaries, monitored and retired securely.
- Initiates reassessment when a material change occurs.
- Reports incidents, complaints and boundary breaches promptly.

Engineering Staff:

- Review agent action logs daily.
- Halt and escalate any agent that acts outside its boundary.
- Review outbound correspondence before sending.
- Report anomalies, drift and supplier notifications to the system owner.

All Personnel:

- Comply with this policy and with the AI Acceptable Use Policy.
- Report suspected AI incidents, misuse or boundary breaches through established channels.

6. AI System Inventory

6.1 The Head of Engineering maintains a central inventory of all AI systems within scope.

6.2 Each inventory entry records at minimum:

- System name and unique identifier.
- System owner.
- Business purpose and intended use.
- Data inputs, including any personal information.
- Data outputs and any actions the system may take.
- Risk tier and date of most recent assessment.
- Supplier or internal development team.
- Hosting location and confirmation that data remains within Canada.
- Action boundary, where applicable.
- Human oversight arrangements.
- Date of last review and next scheduled review.

6.3 The inventory is updated within five business days of any material change, new deployment or retirement.

6.4 The inventory is reviewed in full on a cycle set by Solstice Grid Services; annual review is offered as a starting point and is not derived from a cited control. The review confirms that entries are accurate, that retired systems have been removed, and that no system is operating without an entry.

7. Risk Classification

7.1 Every AI system is assigned a risk tier—High, Medium or Low—based on its potential impact if it fails, behaves unexpectedly or is misused.

7.2 Impact is assessed on two axes, held separately and never averaged:

- Grid safety and reliability.
- Impact on customers or employees, including personal information.

The higher of the two governs the risk tier.

7.3 Likelihood is considered alongside impact. A low-likelihood, high-impact scenario may still warrant a High tier where the consequences are severe.

7.4 Any risk arising from an agent acting without human confirmation of each action is rated High until the action boundary that contains it has been tested and the test results documented.

7.5 Indicative tier assignments:

High:

- Autonomous agents that act on grid telemetry.
- Any system that could affect grid safety, protection settings or load switching if boundaries fail.
- Any system processing personal information for a large population where a failure could cause significant harm.

Medium:

- Systems that draft customer correspondence for human review before sending.
- Systems that inform but do not execute operational decisions.
- Systems processing personal information with moderate impact if compromised.

Low:

- Internal productivity tools with no access to personal information or operational systems.
- Systems whose failure would cause inconvenience but no safety, reliability or privacy harm.

7.6 The system owner proposes the tier. The Head of Engineering confirms Medium and Low assignments. The Chief Technology Officer confirms High assignments.

7.7 The tier is reassessed whenever a material change occurs or at the scheduled review, whichever is sooner.

8. Impact Assessment and Review Before Deployment

8.1 No AI system may be deployed without a documented impact assessment reviewed by the appropriate authority.

8.2 The assessment addresses:

- The business purpose and intended use.
- The data the system will receive, including any personal information.
- The outputs or actions the system will produce.
- The potential harms if the system fails, drifts, is misused or is compromised.
- The controls that reduce those harms.
- The residual risk after controls.
- The human oversight arrangements.
- For agents, the action boundary and the test evidence that the boundary holds.
- Supplier arrangements and compliance with required contract terms.
- Confirmation that no customer data will be used for vendor model training and that data will not leave Canada.

8.3 Assessment depth is proportionate to risk tier:

- High-tier systems require a full impact assessment with documented test evidence for any action boundary.
- Medium-tier systems require a standard assessment covering all elements above.
- Low-tier systems require a brief assessment confirming purpose, data, controls and ownership.

8.4 Review authority:

- High-tier systems: Chief Technology Officer.
- Medium-tier systems: Head of Engineering.
- Low-tier systems: System owner, with the Head of Engineering notified.

8.5 Critical suppliers—the agent orchestration platform, the model provider, the telemetry pipeline vendor and the mapping data provider—require Chief Technology Officer review and completion of the Third-Party AI Review Procedure before any action capability is granted.

8.6 The review record, including the decision and any conditions, is retained with the inventory entry.

9. Material Change

9.1 A material change to an AI system triggers reassessment before the change takes effect.

9.2 Material changes include but are not limited to:

- A new model version from a supplier.
- An expanded or altered action boundary for an agent.
- A change in the data the system receives or the outputs it produces.
- A new subprocessor or a change in hosting jurisdiction.
- A change in the population affected or the sensitivity of data processed.
- A supplier notification of capability change, as required by contract.

9.3 The system owner identifies material changes and initiates reassessment. Where a supplier notifies Solstice Grid Services of a change, the system owner assesses materiality within five business days.

9.4 The reassessment follows the same process and authorities as initial deployment. If the change elevates the risk tier, the higher tier's review authority applies.

9.5 Until reassessment is complete and reviewed, the change does not proceed. If a supplier change cannot be delayed, the system owner escalates to the Chief Technology Officer, who may accept interim operation with additional controls or direct that the system be suspended.

10. Human Oversight and Decision Rights

10.1 Every AI system has assigned human oversight proportionate to its potential impact.

10.2 For autonomous agents that act without human confirmation of each action:

- The agent operates only within a defined action boundary.
- The boundary specifies what the agent may do and what it may not do. Agents may not switch load, may not touch protection settings and may not act on the operational technology network.
- Every action is logged with its inputs and is reversible.
- An engineer reviews the action log daily.
- Any action outside the boundary halts the agent immediately until a human clears it.
- The boundary is tested before deployment and the test evidence is documented.

10.3 For systems that draft customer correspondence:

- Correspondence is reviewed by a human before sending.

- The reviewer confirms accuracy, tone and compliance with applicable requirements.

10.4 Escalation:

- Boundary breaches, unexpected outputs and incidents are escalated to the system owner immediately.
- The system owner escalates to the Head of Engineering where the issue cannot be resolved or where it affects multiple systems.
- The Head of Engineering escalates to the Chief Technology Officer where the issue involves High-tier risk, grid safety or regulatory exposure.

10.5 Decision rights for accepting residual risk:

- High residual risk: Chief Technology Officer only, for a fixed period, with a recorded review date.
- Medium residual risk: Head of Engineering.
- Low residual risk: System owner.

11. Monitoring and Ongoing Assurance

11.1 AI systems are monitored throughout their operational life. Monitoring intensity is proportionate to risk tier.

11.2 Monitoring covers:

- Performance against intended purpose.
- Drift in outputs, accuracy or behaviour.
- Incidents, errors and near-misses.
- Complaints from customers, employees or other affected parties.
- Supplier notifications of model changes, capability changes, incidents or subprocessor changes.
- Continued compliance with action boundaries for agents.
- Continued compliance with data handling requirements, including the prohibition on vendor model training and the requirement that data remain within Canada.

11.3 For High-tier systems, the system owner reports monitoring findings to the Head of Engineering on a cycle set by Solstice Grid Services; monthly reporting is offered as a starting point and is not derived from a cited control.

11.4 Incidents involving AI systems are managed under the AI Incident Response Procedure. Incidents involving personal information are also managed under the Privacy Breach Response Procedure. Security incidents follow the Security Incident Response Procedure.

11.5 Complaints related to AI system outputs or actions are logged, investigated and reported to the system owner. Patterns of complaints trigger reassessment.

11.6 Supplier changes notified under contract are assessed for materiality within five business days. The right to revoke action capability immediately without terminating the contract is exercised where a supplier change introduces unacceptable risk.

12. Exceptions

12.1 Exceptions to this policy are permitted only where a legitimate business need exists and compensating controls reduce risk to an acceptable level.

12.2 Every exception is:

- Documented, stating the requirement excepted, the justification, the compensating controls and the residual risk.
- Time-limited, with a recorded expiry date.
- Accepted by the authority for the relevant risk tier: Chief Technology Officer for High, Head of Engineering for Medium, system owner for Low.
- Reviewed before expiry. An exception that is not reviewed and either renewed or closed by its expiry date is escalated to the Chief Technology Officer.

12.3 The Head of Engineering maintains a register of active exceptions and reports the count of exceptions past their expiry date as a policy metric.

13. Retirement and Evidence Preservation

13.1 When an AI system is no longer required, the system owner initiates retirement.

13.2 Retirement includes:

- Removing the system from operation and revoking any action capabilities.
- Confirming that data processed by the system is handled in accordance with the Records Retention and Disposal Policy and the Privacy and Personal Information Policy.
- Preserving evidence required for regulatory, legal or audit purposes, including impact assessments, review records, action logs and incident records, for the retention periods mandated by the Records Retention and Disposal Policy.
- Updating the inventory to reflect retirement.
- Notifying affected parties where appropriate.

13.3 For systems provided by third parties, retirement includes confirming that the supplier has deleted or returned Solstice Grid Services data in accordance with contract terms.

13.4 The retirement record is retained with the inventory entry.

14. Metrics

The Head of Engineering reports the following metrics to the Chief Technology Officer on a cycle aligned with the policy review frequency:

- Percentage of AI systems with a current risk assessment (assessment completed within the scheduled review period).
- Percentage of High-impact AI systems with assigned and documented human oversight.
- Percentage of required reviews (initial and periodic) completed on time.
- Number of active exceptions past their expiry date.
- Number of material control failures (boundary breaches, incidents, significant drift) and count of overdue corrective actions.

Thresholds for acceptable performance on each metric are set by Solstice Grid Services and recorded in the AI Risk Management Standard; this policy does not prescribe them.

15. Related Documents

This policy operates alongside and should be read with:

- AI Acceptable Use Policy
- AI Risk Management Standard
- AI Incident Response Procedure
- Third-Party AI Review Procedure
- Privacy and Personal Information Policy
- Privacy Breach Response Procedure
- Security Incident Response Procedure
- Information Security Policy
- Data Classification and Handling Standard
- Third-Party Risk Management Policy
- Records Retention and Disposal Policy
- Business Continuity Policy
- Access Control Policy

16. Review and Maintenance

16.1 The Head of Engineering owns this policy and is accountable for its maintenance.

16.2 This policy is reviewed annually or sooner if triggered by:

- A significant AI-related incident.
- A material change in Solstice Grid Services' use of AI.
- A change in applicable law or regulatory guidance.
- Findings from audit or assurance activity.

16.3 Amendments are reviewed by the Chief Technology Officer before taking effect.

16.4 The current version, review history and next scheduled review date are recorded in the document control system.

17. Compliance and Enforcement

17.1 All personnel within scope are expected to comply with this policy.

17.2 Non-compliance may result in disciplinary action, up to and including termination of employment or contract.

17.3 Suspected violations are reported through established channels and investigated promptly.

17.4 The Head of Engineering may direct that an AI system be suspended where continued operation poses unacceptable risk or where the system owner has not remediated identified deficiencies within a reasonable period.

Owner review items

The drafting process identified the following in the supplied answers. Resolve each before customer release.

Type	Detail	Where
gap	No verified framework text was supplied. This policy is drafted against RiskWoven AI practical best practices as stated, but no control identifiers or clause references could be cited because no framework content was supplied.	Entire document
gap	The customer's answers do not specify the retention periods for AI-related records (impact assessments, action logs, review records). The policy defers to the Records Retention and Disposal Policy, but Solstice Grid Services must confirm that policy contains appropriate periods for AI governance evidence.	Section 13 (Retirement and Evidence Preservation)
gap	Metric thresholds (e.g., acceptable percentage of on-time reviews, maximum tolerable overdue exceptions) are not specified in the customer's answers. The policy directs that these be set in the AI Risk Management Standard.	Section 14 (Metrics)
gap	The customer's answers reference applicable Ontario privacy law and the requirement that data not leave Canada, but do not name the specific statutes. Solstice Grid Services' legal advisors should confirm which of PIPEDA, Ontario's privacy legislation or sector-specific regulation applies and whether the Privacy and Personal Information Policy reflects those requirements.	Sections 8, 11, 13
vague	The frequency for full inventory review and for High-tier monitoring reports is stated as an organisational parameter with a recommended starting point. Solstice Grid Services must confirm and document the chosen frequencies in the AI Risk Management Standard.	Sections 6.4 and 11.3

Information the organization must supply

These values were not available from the intake and are required before this document can operate:

- Confirm the retention periods in the Records Retention and Disposal Policy cover AI governance evidence: impact assessments, review records, action logs, incident records and retirement records.
- Set and document metric thresholds in the AI Risk Management Standard (e.g., target percentage for on-time reviews, maximum tolerable overdue exceptions).
- Confirm the frequency for full AI system inventory review (annual is offered as a starting point).
- Confirm the frequency for High-tier system monitoring reports to the Head of Engineering (monthly is offered as a starting point).
- Confirm with legal advisors which privacy statutes apply (PIPEDA, Ontario legislation, sector regulation) and ensure the Privacy and Personal Information Policy reflects them.
- Confirm that the AI Risk Management Standard and Third-Party AI Review Procedure align with the review authorities and supplier requirements stated in this policy.

Written and reviewed by RiskWoven.com. Written by RiskWoven.com's AI drafting agent and checked by RiskWoven.com's independent review stage before release. The review is automated. It is not a human, legal or audit review.